Novell
Open Enterprise Server

www.novell.com

February 1, 2007

NOVELL CLUSTER SERVICES™ 1.8.1
RESOURCE CONFIGURATION GUIDE
FOR LINUX*

Novell®

Legal Notices

Novell, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

You may not use, export, or re-export this product in violation of any applicable laws or regulations including, without limitation, U.S. export regulations or the laws of the country in which you reside.

Copyright © 2005 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at http://www.novell.com/company/legal/patents/ and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

eDirectory is a trademark of Novell, Inc.
NetWare is a registered trademark of Novell, Inc. in the United States and other countries.
NetWare Core Protocol and NCP are trademarks of Novell, Inc.
Novell is a registered trademark of Novell, Inc. in the United States and other countries.
Novell Cluster Services is a trademark of Novell, Inc.
Novell Directory Services and NDS are registered trademarks of Novell, Inc. in the United States and other countries.
Novell Storage Services is a trademark of Novell, Inc.
SUSE is a registered trademark of SUSE LINUX AG, a Novell business.

Third-Party Materials

All third-party trademarks are the property of their respective owners.

Contents

About This Guide 7

1 Configuring DHCP with Novell Cluster Services 9

1.1 Benefits of Configuring DHCP for High Availability . . . 9
1.2 DHCP Installation and Configuration . . . 9
  1.2.1 Configuring DHCP on the Shared Disk . . . 9
  1.2.2 Configuring the dhcpd.conf File . . . 10
  1.2.3 Creating a dhcpd.leases File . . . 10
1.3 Novell Cluster Services Configuration and Setup . . . 10
  1.3.1 Creating a DHCP Cluster Resource . . . 10
  1.3.2 Configuring DHCP Load and Unload Scripts . . . 11
  1.3.3 Setting DHCP Start, Failover, and Failback Modes . . . 12
  1.3.4 View or Edit DHCP Resource Server Assignments . . . 13
1.4 Additional Information . . . 13

2 Configuring DNS with Novell Cluster Services 15

2.1 Benefits of Configuring DNS for High Availability . . . 15
2.2 DNS Installation and Configuration . . . 15
  2.2.1 Configuring DNS on the Shared Disk . . . 15
  2.2.2 Configuring the named.conf File . . . 16
  2.2.3 Creating Zone Files . . . 16
2.3 Novell Cluster Services Configuration and Setup . . . 16
  2.3.1 Creating a DNS Cluster Resource . . . 17
  2.3.2 Configuring DNS Load and Unload Scripts . . . 17
  2.3.3 Setting DNS Start, Failover, and Failback Modes . . . 18
  2.3.4 View or Edit DNS Resource Server Assignments . . . 19
2.4 Additional Information . . . 19

3 Configuring iPrint with Novell Cluster Services 21

3.1 Understanding the Benefits of High Availability Printing . . . 21
3.2 iPrint Cluster Requirements . . . 22
3.3 OES Server Installation Options . . . 22
3.4 Configuring a Shared Disk . . . 22
3.5 Carving Disk Partitions . . . 23
3.6 Update Your Server . . . 23
3.7 Use Cfdisk to Prepare the Partition . . . 23
3.8 Configuring Novell Cluster Services . . . 23
3.9 Creating Traditional Linux File Systems on Shared Disks . . . 24
  3.9.1 Ensuring That the Shared Disk Is not a Compatibility Volume . . . 25
  3.9.2 Removing Other Segment Managers . . . 27
  3.9.3 Creating a Cluster Segment Manager Container . . . 27
  3.9.4 Make New Container Active . . . 28
  3.9.5 Creating an EVMS Volume for iPrint Resources (Conditional) . . . 29
  3.9.6 Creating a File System on the EVMS Volume . . . 29
3.10 Configure a Cluster Resource to Host Printing . . . 30
3.11 Starting iPrint Cluster Resource . . . 32
3.12 Creating and Configuring a Driver Store and Print Manager . . . 33

3.13 Changing Start, Failover, and Failback Modes for Printing . . . 34
3.14 Using iPrint in a Clustered Environment . . . 34
3.15 Troubleshooting iPrint in a Cluster . . . 35
  3.15.1 Error Code 506D0203 . . . 35
  3.15.2 Resource Goes Comatose on Failover . . . 35
  3.15.3 iPrint Web Pages Are Displayed on One Node Only . . . 35
3.16 iPrint Configuration Worksheet . . . 35

4 Configuring Samba with Novell Cluster Services 37

4.1 Benefits of Configuring Samba for High Availability . . . 37
4.2 Samba Installation and Configuration . . . 37
  4.2.1 Configuring Samba on the Shared Disk . . . 37
  4.2.2 Configuring the smb.conf File . . . 38
4.3 Novell Cluster Services Configuration and Setup . . . 38
  4.3.1 Creating a Samba Cluster Resource . . . 39
  4.3.2 Configuring Samba Load and Unload Scripts . . . 39
  4.3.3 Setting Samba Start, Failover, and Failback Modes . . . 41
  4.3.4 View or Edit Samba Resource Server Assignments . . . 42
4.4 Additional Information . . . 42

About This Guide

Before running any application on servers in a Novell® Cluster Services™ cluster, the application must be installed and configured properly to take advantage of shared storage and to achieve high availability. This document provides information to help you set up and configure some of the applications and services commonly used in an OES Linux environment. See the Novell Cluster Administration Guide for Linux for cluster installation and configuration instructions, and see your application documentation for specific application installation and setup instructions.

Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (®, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as UNIX(thrdmrk) or Linux*, should use forward slashes as required by your software.

User Comments

We want to hear your comments and suggestions about this manual and the other documentation included with OES. To contact us, use the User Comments feature at the bottom of any page in the online documentation.

Configuring DHCP with Novell Cluster Services

1

1.1 Benefits of Configuring DHCP for High Availability

A Novell® OES Linux DHCP server automatically assigns IP addresses and other configuration information to clients upon request or when the clients are restarted. If for some reason the Novell OES Linux DHCP server is not accessible, clients lose their ability to connect to the network because they cannot obtain an IP address. Configuring DHCP with Novell Cluster Services™ ensures the IP address range required by users to connect to the network is highly available. This is possible because the DHCP server is automatically started, stopped, and restarted on different servers in the cluster by Novell Cluster Services.

1.2 DHCP Installation and Configuration

Novell Cluster Services (NCS) 1.8 must be installed before configuring DHCP to work in a cluster. NCS 1.8 provides a DHCP resource template, which facilitates configuring DHCP in a cluster environment. Novell OES includes NCS 1.8 and licenses for two cluster nodes.

Before running DHCP with Novell Cluster Services, DHCP must be installed on every server that will run it. To do this, install the Internet Software Consortium (ISC) DHCP server package that ships with SLES9. This server package can be installed when you install your SLES9 server or afterward. You can install the server package after the SLES9 installation using either YaST or Linux commands. More information on DHCP configuration is available in the /usr/share/doc/packages/dhcp-server directory.

1.2.1 Configuring DHCP on the Shared Disk

To configure DHCP to use shared storage, you need to create a shared directory (file system or disk) on the shared disk system and create mount points to that shared file system on each cluster server that will run DHCP.

1. Create a file system on the shared disk system using EVMS.

   Enter `evmsgui` at the Linux server to start the EVMSGUI utility.

   Click Help and go the the Web links provided for documentation and information on using EVMS to configure and manage storage.
2. On each cluster node that will run DHCP, create the mount point.

   For example, the mount point could be /mnt/dhcp.
3. Log in as the root user and mount the shared disk (file system) that was created in Step 1.

   For example, depending on the mount point and directory names, you could enter a command similar to the following to mount the shared disk:

```
mount /dev/evms/dhcp /mnt/dhcp
```

4. At the root of the shared disk you just created (/mnt/dhcp), enter the following commands to create the directories specified:

   ```
   mkdir etc
   ```

   ```
   mkdir -p var/lib/dhcp
   ```

   These directories must be owned by the root user, and the default group must be root. Also, all four directories must have permissions of drwxr-xr-x.

1.2.2 Configuring the dhcpd.conf File

1. Copy the dhcpd.conf file from the /etc directory on one of the OES Linux cluster servers to the etc directory you created on the shared disk in Step 4 on page 10.

   This would be the /mnt/dhcp/etc directory if you used the same directory names as those given in the example above.

2. Configure the DHCP daemon by modifying the dhcpd.conf that you copied in Step 1.

   For information on what to modify, see any of the following:

   man dhcpd.conf
   http://www.isc.org (http://www.isc.org)
   The Linux Documentation Project http://www.tldp.org (http://www.tldp.org)

1.2.3 Creating a dhcpd.leases File

A dhcpd.leases file is necessary for DHCP to function. The DHCP daemon requires this file before it will start. The file can be empty, and it must reside in the var/lib/dhcp directory you created in Step 4 on page 10.

One way to create the empty file is to use the touch command. For example, if you used the directory names listed in the example above, you could enter the following to create an empty dhcpd.leases file:

```
touch /mnt/dhcp/var/lib/dhcp/dhcpd.leases
```

1.3 Novell Cluster Services Configuration and Setup

After DHCP is properly installed and configured , you must create and configure a DHCP resource in Novell Cluster Services. This includes configuring DHCP load and unload scripts; setting DHCP start, failover, and failback modes; and assigning the DHCP resource to specific servers in your cluster.

1.3.1 Creating a DHCP Cluster Resource

Novell Cluster Services includes a DHCP resource template, which greatly simplifies the process for creating a DHCP cluster resource. Much of the DHCP cluster resource configuration is perfomed automatically by the DHCP resource template. See "Creating Cluster Resources" in the *OES Novell Cluster Services 1.8.1 Administration Guide for Linux* for more information.

To create a DHCP cluster resource:

1. Ensure the shared disk (file system) you created and mounted in Step 2 on page 9 is unmounted.

   If you used the directory names specified in the example, you can enter `umount /mnt/dhcp` to unmount the shared disk.
2. Open your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/imanager.html. Replace *server_ip_address* with the IP address or DNS name of a server in the cluster or with the IP address for Apache-based services.
3. Enter your username and password.
4. In the left column, locate Clusters and then click the Cluster Options link.

   iManager displays three links under Clusters that you can use to configure and manage your cluster.
5. Specify the cluster name, or browse and select it, then click the New link.
6. Specify Resource as the resource type you want to create by clicking the Resource radio button, then click Next.
7. Specify a name for the DHCP resource.

   NOTE: Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space will be converted to an underscore.
8. Type the DHCP template name in the Inherit From Template field, or browse and select it from the list.
9. Check the Define Additional Properties check box, click Next, and then continue with "Configuring DHCP Load and Unload Scripts" on page 11.

   The DHCP resource template configures the DHCP resource by automatically creating DHCP load and unload scripts, setting failover and failback modes, and assigning DHCP as a resource to all nodes in the cluster.

1.3.2 Configuring DHCP Load and Unload Scripts

Load Script Configuration

The DHCP load script page should already be displayed. The load script contains commands to start the DHCP service. You must customize some commands for your specific DHCP configuration.

The load script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# Template for clustering the ISC DHCP daemon.
# You should not configure failover within the DHCP server
# itselft. Instead, let NCS manage the failover.
# mount the file system.
exit_on_error mount -t reiserfs /dev/evms/vol /mnt/dhcp
# add the IP address
```

```
exit_on_error add_secondary_ipaddress a.b.c.d
# start dhcpd
exit_on_error /usr/sbin/dhcpd -cf /mnt/dhcp/etc/dhcpd.conf -lf /mnt/
dhcp/var/lib/dhcp/dhcpd.leases
# return status
exit 0
```

To Customize the DHCP load script for your specific configuration:

1. View and, if necessary, edit the following lines for your specific directory structure, mount point, and file system.

   ```
   exit_on_error mount -t reiserfs /dev/evms/vol /mnt/dhcp
   exit_on_error /usr/sbin/dhcpd -cf /mnt/dhcp/etc/dhcpd.conf -lf /
   mnt/dhcp/var/lib/dhcp/dhcpd.leases
   ```

2. Edit the following line to assign a unique IP address to the DHCP cluster resource.

   ```
   exit_on_error add_secondary_ipaddress a.b.c.d
   ```

   Replace a.b.c.d with the IP address you want to assign to the DHCP cluster resource.

   The IP address for the DHCP cluster resource allows clients to reconnect to that address regardless of which server is hosting it.

3. Click Next and continue with the unload script configuration.

Unload Script Configuration

The DHCP unload script page should now be displayed. The unload script contains commands to stop the DHCP service. You must customize some commands for your specific DHCP configuration.

The unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# request dhcpd stop
ignore_error killproc -p /var/run/dhcpd.pid -TERM /usr/sbin/dhcpd
# stop dhcpd otherwise
sleep 8
ignore_error fuser -k /mnt/dhcp
# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d
# umount the file system
exit_on_error umount /mnt/dhcp
# return status
exit 0
```

1. View and, if necessary edit the mount points in the applicable lines of the unload script.
2. Replace a.b.c.d with the same IP address you specified in the load script, then click Next.

1.3.3 Setting DHCP Start, Failover, and Failback Modes

The page to set Start, Failover, and Failback modes should now be displayed.

The DHCP resource template sets the DHCP resource Start mode and Failover mode to Auto and the Failback Mode to Disable.

If the DHCP resource Start mode is set to Auto, the DHCP server automatically loads on a designated server when the cluster is first brought up. If the DHCP Start mode is set to Manual, you can manually start DHCP on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the DHCP Failover mode is set to Auto, the DHCP server automatically moves to the next server in the Assigned Nodes list in the event of a hardware or software failure. If the DHCP Failover mode is set to Manual, you can intervene after a failure occurs and before the DHCP server is started on another node.

If the DHCP Failback mode is set to Disable, the DHCP server will continue running on the node it has failed to. If the DHCP Failback Mode is set to Auto, the DHCP server automatically moves back to its preferred node when the preferred node is brought back online. Set the DHCP Failback mode to Manual to prevent the DHCP server from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

View or change the DHCP resource Start, Failover, and Failback modes, then click Next and continue with "View or Edit DHCP Resource Server Assignments" on page 13.

1.3.4 View or Edit DHCP Resource Server Assignments

The page to view or change DHCP resource server assignments should now be displayed. The DHCP resource template automatically assigns the DHCP resource to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list.

To view or edit DHCP resource node assignments or change the server failover order:

1. From the list of unassigned nodes, select the server you want the resource assigned to, then click the Right-arrow button to move the selected server to the Assigned Nodes list.

   Repeat this step for all servers you want assigned to the resource. You can also use the Left-arrow button to unassign servers from the resource.
2. Click the Up-arrow and Down-arrow buttons to change the failover order of the servers assigned to the resource or volume.
3. Click Apply or Finish to save node assignment changes.

1.4 Additional Information

The configuration described in this document does not run the DHCP daemon in a chroot jail. If that level of security is required, additional configuration is necessary.

For additional information on configuring the ISC DHCP daemon, see the following:

man/info dhcpd
man/info dhcpd.conf
man/info dhcpd.leases
http://www.isc.org (http://www.isc.org)
The Linux Documentation Project (http://www.tldp.org)

Configuring DNS with Novell Cluster Services

2

A Novell® OES Linux DNS server automatically resolves IP addresses with domain names. If for some reason the Novell OES Linux DNS server is not accessible, clients will lose their ability to access certain network resources and applications. Configuring DNS with Novell® Cluster Services™ can help avoid this loss of access.

2.1 Benefits of Configuring DNS for High Availability

Configuring DNS with Novell Cluster Services ensures the IP address/domain name resolution required by users is always highly available. This is possible because the DNS server is automatically started, stopped, and restarted on different servers in the cluster by Novell Cluster Services.

2.2 DNS Installation and Configuration

Novell Cluster Services (NCS) 1.8 must be installed before configuring DNS to work in a cluster. NCS 1.8 provides a DNS resource template, which facilitates configuring DNS in a cluster environment. Novell OES includes NCS 1.8 and licenses for two cluster nodes.

Before running DNS with Novell Cluster Services, DNS must be installed on every server that will run it. To do this, install the DNS DHCP server package that ships with SLES 9. This server package can be installed when you install your SLES 9 server or afterward. You can install the server package after the SLES 9 installation by installing the Bind Domain Name Server component using either YaST or Linux commands. More information on DNS configuration is available in the /usr/share/doc/packages/bind directory.

2.2.1 Configuring DNS on the Shared Disk

To configure DNS to use shared storage, you need to create a shared directory (file system or disk) on the shared disk system and create mount points to that shared file system on each cluster server that will run DNS.

1. Create a file system on the shared disk system using EVMS.

   Enter `evmsgui` at the Linux server to start the EVMSGUI utility.

   Click Help and use the Web links provided for documentation and information on using EVMS to configure and manage storage.
2. On each cluster node that will run DNS, create the mount point.

   For example, the mount point could be /mnt/named.
3. At the root of the shared disk you just created (/mnt/named), enter the following commands to create the directories specified:

   ```
   mkdir etc
   ```

```
mkdir -p var/lib/named
```

These directories must be owned by the root user, and the default group must be root. Also, all four directories must have permissions of drwxr-xr-x.

4. Log in as the root user and mount the shared disk (file system) that was created in Step 1.

   For example, depending on the mount point and directory names, you could enter a command similar to the following to mount the shared disk:

   ```
   mount /dev/evms/named /mnt/named
   ```

5. Copy the root.hint file from the /var/lib/named directory on one of the OES Linux cluster servers to the /var/lib/named directory on the shared disk.

2.2.2 Configuring the named.conf File

1. Copy the named.conf file from the /etc directory on one of the OES Linux cluster servers to the etc directory you created on the shared disk in Step 4 on page 10.

   This would be the /mnt/named/etc directory if you used the same directory names as those given in the example above.

2. Configure the DNS daemon by modifying the named.conf that you copied in Step 1.

   For information on what to modify, see any of the following:

   man/info named.conf
   http://www.isc.org (http://www.isc.org)
   The Linux Documentation Project http://www.tldp.org (http://www.tldp.org)
   The BIND 9 Administrator Reference Manual

2.2.3 Creating Zone Files

Create the necessary zone files in the /var/lib/named directory you created in Step 4 on page 10.

This would be the /mnt/named/var/lib/named directory if you used the same directory names as those give in the example in Step 2 on page 9.

For more information, see any of the following:

http://www.isc.org (http://www.isc.org)
The Linux Documentation Project http://www.tldp.org (http://www.tldp.org)
The BIND 9 Administrator Reference Manual

2.3 Novell Cluster Services Configuration and Setup

After DNS is properly installed and configured , you must create and configure a DNS resource in Novell Cluster Services. This includes configuring DNS load and unload scripts; setting DNS Start, Failover, and Failback modes; and assigning the DNS resource to specific servers in your cluster.

2.3.1 Creating a DNS Cluster Resource

Novell Cluster Services includes a DNS resource template, which greatly simplifies the process for creating a DNS cluster resource. Much of the DNS cluster resource configuration is perfomed automatically by the DNS resource template. See "Creating Cluster Resources" in the *OES Novell Cluster Services 1.8.1 Administration Guide for Linux* for more information.

To create a DNS cluster resource:

1. Ensure the shared disk (file system) you created and mounted in Step 1 on page 9 is unmounted.

   If you used the directory names specified in the example, you can enter `umount /mnt/named` to unmount the shared disk.
2. Open your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/imanager.html. Replace *server_ip_address* with the IP address or DNS name of a server in the cluster or with the IP address for Apache-based services.
3. Enter your username and password.
4. In the left column, locate Clusters and then click Cluster Options.

   iManager displays three links under Clusters that you can use to configure and manage your cluster.
5. Specify the cluster name, or browse and select it, then click New.
6. Specify Resource as the resource type you want to create by clicking the Resource radio button, then click Next.
7. Specify a name for the DNS resource.
8. Type the DNS template name in the Inherit From Template field, or browse and select it from the list.
9. Check the Define Additional Properties check box, click Next, and then continue with "Configuring DHCP Load and Unload Scripts" on page 11.

   The DNS resource template configures the DNS resource by automatically creating DNS load and unload scripts, setting failover and failback modes, and assigning DNS as a resource to all nodes in the cluster.

2.3.2 Configuring DNS Load and Unload Scripts

Load Script Configuration

The DNS load script page should already be displayed. The load script contains commands to start the DNS service. You must customize some commands for your specific DNS configuration.

The load script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# mount the file system.
exit_on_error mount -t reiserfs /dev/evms/named /mnt/named
# add the IP address
exit_on_error add_secondary_ipaddress a.b.c.d
```

```
# start named
exit_on_error /usr/sbin/named -c /mnt/named/etc/named.conf -u named
sleep 2
exit_on_error checkproc -p /var/run/named/named.pid /usr/sbin/named
# return status
exit 0
```

To customize the DNS load script for your specific configuration:

1. View and, if necessary, edit the following lines for your specific directory structure, mount point, and file system:
   ```
   exit_on_error mount -t reiserfs /dev/evms/named /mnt/named
   exit_on_error /usr/sbin/named -c /mnt/named/etc/named.conf -u named
   exit_on_error checkproc -p /var/run/named/named.pid /usr/sbin/
   named
   ```
2. Edit the following line to assign a unique IP address to the DNS cluster resource:
   ```
   exit_on_error add_secondary_ipaddress a.b.c.d
   ```
   Replace a.b.c.d with the IP address you want to assign to the DNS cluster resource.

   The IP address for the DNS cluster resource allows clients to reconnect to that address regardless of which server is hosting it.
3. Click Next and continue with the unload script configuration.

Unload Script Configuration

The DNS unload script page should now be displayed. The unload script contains commands to stop the DNS service. You must customize some commands for your specific DNS configuration.

The unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# request named stop
ignore_error killproc -p /var/run/named.pid -TERM /usr/sbin/named
# stop named otherwise
sleep 8
ignore_error fuser -k /mnt/named
# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d
# umount the file system
exit_on_error umount /mnt/named
# return status
exit 0
```

1. View and, if necessary, edit the mount points in the applicable lines of the unload script.
2. Replace a.b.c.d with the same IP address you specified in the load script, then click Next.

2.3.3 Setting DNS Start, Failover, and Failback Modes

The page to set Start, Failover, and Failback modes should now be displayed.

The DNS resource template sets the DNS resource Start mode and Failover mode to Auto and the Failback Mode to Disable.

If the DNS resource Start mode is set to Auto, the DNS server automatically loads on a designated server when the cluster is first brought up. If the DNS Start mode is set to Manual, you can manually start DNS on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the DNS Failover mode is set to Auto, the DNS server automatically moves to the next server in the Assigned Nodes list in the event of a hardware or software failure. If the DNS Failover mode is set to Manual, you can intervene after a failure occurs and before the DNS server is started on another node.

If the DNS Failback mode is set to Disable, the DNS server will continue running on the node it has failed to. If the DNS Failback Mode is set to Auto, the DNS server automatically moves back to its preferred node when the preferred node is brought back online. Set the DNS Failback mode to Manual to prevent the DNS server from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

View or change the DNS resource Start, Failover, and Failback modes, then click Next and continue with "View or Edit DHCP Resource Server Assignments" on page 13.

2.3.4 View or Edit DNS Resource Server Assignments

The page to view or change DNS resource server assignments should now be displayed. The DNS resource template automatically assigns the DNS resource to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list.

To view or edit DNS resource node assignments or change the server failover order:

1. From the list of unassigned nodes, select the server you want the resource assigned to, then click the Right-arrow button to move the selected server to the Assigned Nodes list.

   Repeat this step for all servers you want assigned to the resource. You can also use the Left-arrow button to unassign servers from the resource.
2. Click the Up-arrow and Down-arrow buttons to change the failover order of the servers assigned to the resource or volume.
3. Click Apply or Finish to save node assignment changes.

2.4 Additional Information

The DNS cluster resource template runs the DNS daemon as the user named. Also, the configuration described in this document does not run the DNS daemon in a chroot jail. If that level of security is required, additional configuration is necessary.

For additional information on configuring the ISC DHCP daemon, see the following:

man/info named
http://www.isc.org (http://www.isc.org)
The Linux Documentation Project (http://www.tldp.org)
The BIND 9 Administrator Reference Manual

Configuring iPrint with Novell Cluster Services

3

This section will help you do the following

- Section 3.1, "Understanding the Benefits of High Availability Printing," on page 21
- Section 3.2, "iPrint Cluster Requirements," on page 22
- Section 3.3, "OES Server Installation Options," on page 22
- Section 3.4, "Configuring a Shared Disk," on page 22
- Section 3.5, "Carving Disk Partitions," on page 23
- Section 3.6, "Update Your Server," on page 23
- Section 3.7, "Use Cfdisk to Prepare the Partition," on page 23
- Section 3.8, "Configuring Novell Cluster Services," on page 23
- Section 3.9, "Creating Traditional Linux File Systems on Shared Disks," on page 24
- Section 3.10, "Configure a Cluster Resource to Host Printing," on page 30
- Section 3.11, "Starting iPrint Cluster Resource," on page 32
- Section 3.12, "Creating and Configuring a Driver Store and Print Manager," on page 33
- Section 3.13, "Changing Start, Failover, and Failback Modes for Printing," on page 34
- Section 3.14, "Using iPrint in a Clustered Environment," on page 34
- Section 3.15, "Troubleshooting iPrint in a Cluster," on page 35
- Section 3.16, "iPrint Configuration Worksheet," on page 35

3.1 Understanding the Benefits of High Availability Printing

Running iPrint on a cluster provides the following benefits:

- Centralized management
  - By configuring the Print Management Database to reside on a shared disk cluster volume, you no longer need to replicate information to achieve high availability.
  - The cluster console GUI lets you monitor and control the location of cluster volumes with the Driver Store and Print Manager services using them from a single management workstation.
  - Because cluster volumes are tightly consistent and highly available, distributed print management tasks are simplified.
- Automatic restart of print services
  - If a server fails, the cluster volume containing the Print Management Database and spool area automatically remounts on a surviving server in the cluster. The Driver Store and Print Manager also automatically restart without user intervention.

- On restart, the Driver Store and Print Manager discover the Print Management Database on a cluster volume mounted on their server. Because the cluster volume is the same regardless of which server it is mounted on, no print-related information is lost or out of date.

3.2 iPrint Cluster Requirements

To properly set up and configure iPrint in a cluster environment, you need the IP addresses and partitionslisted below. See Section 3.16, “iPrint Configuration Worksheet,” on page 35 for a worksheet to help in configuring your system.

- IP address for the cluster.
- IP address for each node in the cluster.
- IP address for the iPrint resource in the cluster.
- SBD partition of 10 MB.
- Resource Partition on your shared disk of 20 GB.

3.3 OES Server Installation Options

When you install OES on the server, you should select the iPrint, Novell Cluster Services and NSS options. On the Installation Settings page which summarizes the choices you have made, you must click Novell Cluster Services before you are allowed to continue. Confirm the Admin user's username and password, and then on the next screen you must select *Configure Later*

3.4 Configuring a Shared Disk

After you have installed Novell Cluster Services, you can begin configuring the shared storage device for the cluster as specified by the shared storage device's documentation.

If YAST appears with a new disk found message during the configuration of a iSCSI or a Storage Area Network (SAN), close the dialog. You do not want to create a file system at this point.

WARNING: iPrint is not supported on NSS Pools. You must choose a native Linux file system such as Reiser or EXT3.

SAN: Is a high-speed subnetwork of shared storage devices used for storing data.

For information on configuring a SAN, consult the SAN vendor's documentation.

iSCSI: Uses SCSI block storage protocols networked over a high-speed TCP/IP network to create a low-cost Storage Area Network (SAN) using commodity high-speed Ethernet hardware.

- **iSCSI for NetWare** lets you configure a NetWare server to serve as an iSCSI target. This requires you to install an iSCSI driver that acts as an iSCSI protocol initiator. For more information on iSCSI for NetWare, see *iSCSI 1.1.3 Administration Guide for NetWare 6.5*. For information on the using an iSCSI initiator on Linux, see “Accessing iSCSI Targets on NetWare Servers from Linux Initiators” in the *iSCSI 1.1.3 Administration Guide for NetWare 6.5*.
- **Third-party** vendors have their own iSCSI solutions. Refer to the vendor's documentation.

IMPORTANT: Configure the iSCSI drives to load when the servers starts by entering `chkconfig iscsi on` at a terminal prompt.

3.5 Carving Disk Partitions

Using utilities that came with your shared disk, carve two partitions as follows:

- SBD partition of 10 MB.

  WARNING: Do not reconfigure or remove a logical disk or segment manager containing the sbd partition.

- Resource Partition on your shared disk of 20 GB.

3.6 Update Your Server

These procedures require that you install the listed updates using the OES Red Carpet channel on each node of the cluster. You must be running the listed rpm version or later. See "Patching an OES Linux Server" in the *OES Linux Installation Guide* for more information.

To view the installed rpm versions use the `rpm -qa` command.

novell-ncp-5.0.1-48
novell-ncpserv-1.3-6
novell-nss-4.9.20-1
evms-2.3.3-0.43
novell-cluster-services-km-1.8.2-57

3.7 Use Cfdisk to Prepare the Partition

To ensure your disk is properly partioned and the partition table is zeroed out, use `cfdisk` and complete the following steps.

1. At the command line, enter `cfdisk /dev/sba`.

   An error usually displays indicating the partion is invalid.
2. When prompted, select *Yes* to zero out the partition.
3. Select *Write* from the menu to write out the partition tables.

3.8 Configuring Novell Cluster Services

After you have installed OES, you can begin configuring Novell Cluster Services on all nodes of the cluster.

1. In Yast go to System, and select Cluster Services (NCS).

   Only click once as this might take a moment to initialize.

2 Follow the prompts to configure the service and create an NCS Cluster object.

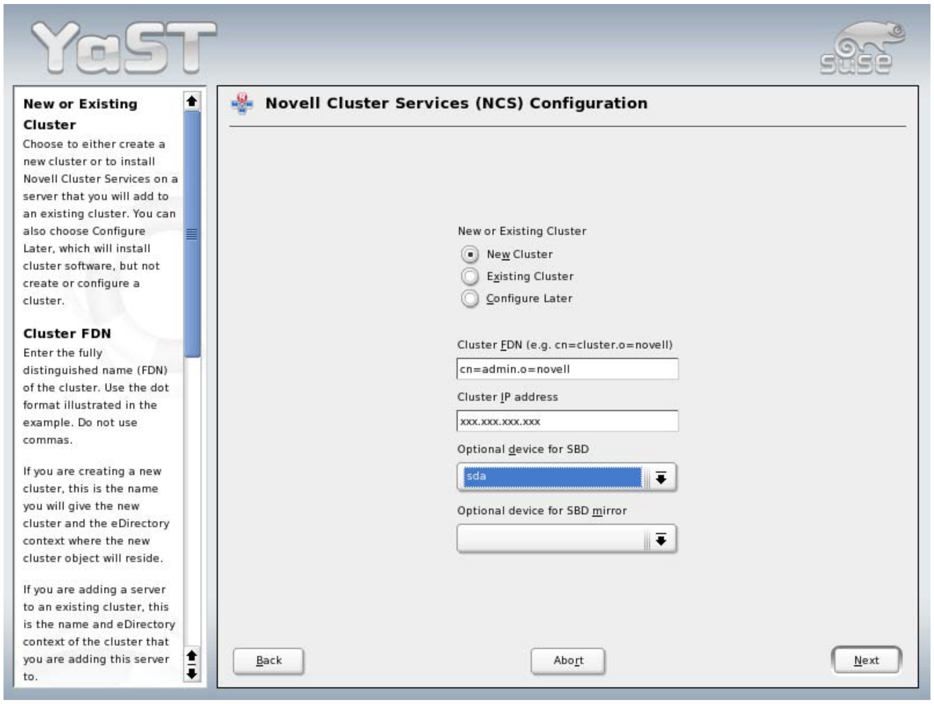

New or Existing Cluster
Choose to either create a new cluster or to install Novell Cluster Services on a server that you will add to an existing cluster. You can also choose Configure Later, which will install cluster software, but not create or configure a cluster.

Cluster FDN
Enter the fully distinguished name (FDN) of the cluster. Use the dot format illustrated in the example. Do not use commas.

If you are creating a new cluster, this is the name you will give the new cluster and the eDirectory context where the new cluster object will reside.

If you are adding a server to an existing cluster, this is the name and eDirectory context of the cluster that you are adding this server to.

Novell Cluster Services (NCS) Configuration

New or Existing Cluster
● New Cluster
○ Existing Cluster
○ Configure Later

Cluster FDN (e.g. cn=cluster.o=novell): cn=admin.o=novell

Cluster IP address: xxx.xxx.xxx.xxx

Optional device for SBD: sda

Optional device for SBD mirror:

To create the object you need an IP address for the cluster resource and the device path for the SBD partition that was defined when you configured the shared disk array. If you cannot see the shared device for the SBD partition (the sda or sdb partion created earlier), repeat the steps in Section 3.7, “Use Cfdisk to Prepare the Partition,” on page 23.

You can look in /dev/disk/by-id/ for this information

When configuring subsequent nodes, choose *Existing Cluster*.

3 Click Next

4 Choose the IP Address for the cluster, and then next.

The cluster object is created and started.

NOTE: The Admin user must be a LUM enabled user in order to manage Clustering and NSS from iManager. In some post installs of Clustering and NSS, the Admin user might not be automatically LUM enabled, and you might be required to manually LUM enable the Admin user.

3.9 Creating Traditional Linux File Systems on Shared Disks

To configure shared media for a cluster, you must use EVMS. Using EVMS to create partitions, containers, volumes, and file systems helps prevent data corruption caused by multiple nodes accessing the same data. You can create partitions and volumes using any of the journaled Linux file systems (EXT3, Reiser, etc.).

EVMS virtual volumes are recommended for Novell Cluster Services because they can more easily fail over to different cluster servers and provide greater protection of shared media from unauthorized access. You can enter `man evms` at the Linux server console to reference the evms man page, which provides additional instructions and examples for evms.

You can also enter `man mount` at the Linux server console to reference the mount man page, which provides additional instructions and examples for the `mount` command.

The following sections provide the necessary information for using EVMS to create a traditional Linux volume and file system on a shared disk:

- Section 3.9.1, "Ensuring That the Shared Disk Is not a Compatibility Volume," on page 25
- Section 3.9.2, "Removing Other Segment Managers," on page 27
- Section 3.9.3, "Creating a Cluster Segment Manager Container," on page 27
- Section 3.9.4, "Make New Container Active," on page 28
- Section 3.9.5, "Creating an EVMS Volume for iPrint Resources (Conditional)," on page 29
- Section 3.9.6, "Creating a File System on the EVMS Volume," on page 29

WARNING: EVMS administration utilities (evms, evmsgui, and evmsn) should not be running when they are not being used. EVMS utilities lock the EVMS engine, which prevents other evms-related actions from being performed. This affects both NSS and traditional Linux volume actions.

NSS and traditional Linux volume cluster resources should not be migrated, offlined, or onlined while any of the EVMS administration utilities are running.

3.9.1 Ensuring That the Shared Disk Is not a Compatibility Volume

New EVMS volumes are by default configured as compatibility volumes. If any of the volumes on your shared disk (that you plan to use in your cluster) are compatibility volumes, you must delete them because the volumes are created later

---

NOTE: The SBD partition should not be changed using this procedure. The SBD partition is fine as a NWSegMgr.

---

1 At the Linux server console, enter `evmsgui`.

2 Click Yes, if the following prompt appears.

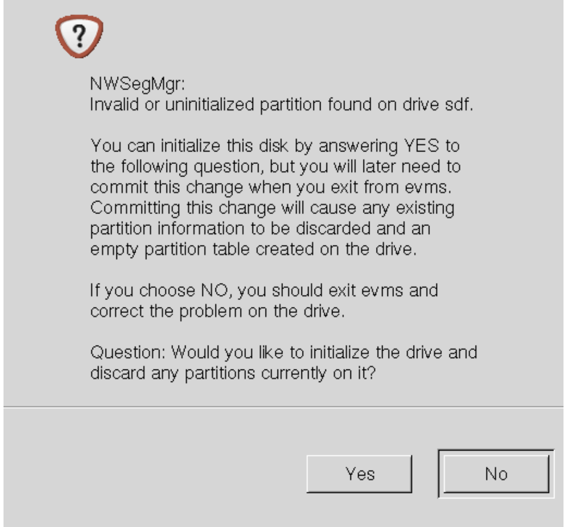

NWSegMgr:
Invalid or uninitialized partition found on drive sdf.

You can initialize this disk by answering YES to the following question, but you will later need to commit this change when you exit from evms. Committing this change will cause any existing partition information to be discarded and an empty partition table created on the drive.

If you choose NO, you should exit evms and correct the problem on the drive.

Question: Would you like to initialize the drive and discard any partitions currently on it?

When starting, EVMS discovers new disks available to the server and configures them as NWSegMgr.

3 Click the *Volumes* tab, then right-click the volume on the shared disk and select *Display details*.

4 Click the *Page 2* tab and determine from the Status field if the volume is a compatibility volume.

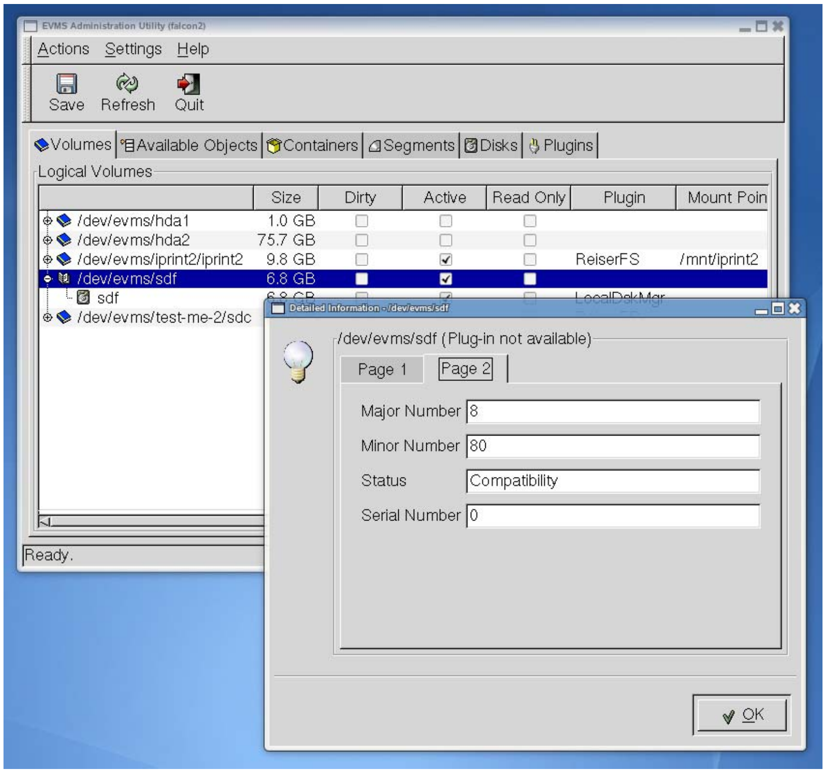

	Size	Dirty	Active	Read Only	Plugin	Mount Poin
/dev/evms/hda1	1.0 GB	☐	☐	☐		
/dev/evms/hda2	75.7 GB	☐	☐	☐		
/dev/evms/iprint2/iprint2	9.8 GB	☐	☑	☐	ReiserFS	/mnt/iprint2
/dev/evms/sdf	6.8 GB	☐	☑	☐		
sdf	6.8 GB	☐	☑	☐	LocalDskMgr	
/dev/evms/test-me-2/sdc						

/dev/evms/sdf (Plug-in not available)

Major Number 8
Minor Number 80
Status Compatibility
Serial Number 0

If the volume is a compatibility volume or has another segment manager on it, continue with Step 4a below.

4a Click the *Volumes* tab, right-click the volume, then select *Delete*.

4b Select the volume, then click *Recursive Delete*.

4c (Conditional) If a Response Required pop-up appears, click the *Write zeros* button.

4d (Conditional) If another pop-up appears, click *Continue* to write 1024 bytes to the end of the volume.

3.9.2 Removing Other Segment Managers

If any of the shared disks you plan to use with your cluster have segment managers other than CSM, you must delete them as well.

The sbd partition cannot be located on the same virtual disk as your cluster resource. If the segment manager contains the sbd partition, use your shared storage tools to carve out a different partition for your cluster resource.

WARNING: Do not reconfigure or remove a logical disk or segment manager containing the sbd partition.

---

1. In evmsgui, click the *Disks* tab, then right-click the disk you plan to use for a cluster resource.
2. Select *remove segment manager from Object*.

   This option only appears if there is another segment manager for the selected disk.

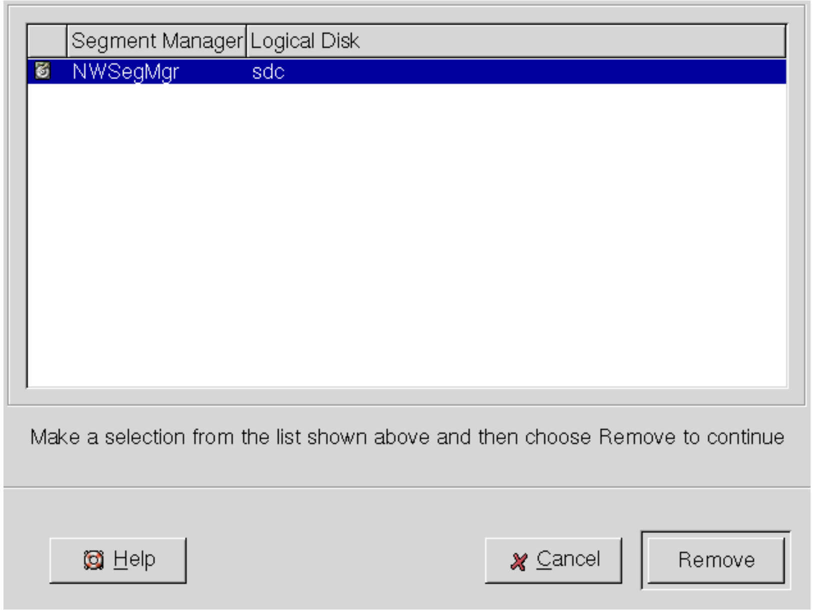

Segment Manager	Logical Disk
NWSegMgr	sdc

Make a selection from the list shown above and then choose Remove to continue

3. Select the listed segment manager and click *Remove*.

3.9.3 Creating a Cluster Segment Manager Container

To use a traditional Linux volume with EVMS as a cluster resource, you must use the Cluster Segment Manager (CSM) plug-in for EVMS to create a CSM container.

You should note the container name used as it is needed when configuring the cluster scripts.

NOTE: CSM containers require Novell Cluster Services (NCS) to be running on all nodes that access the CSM container. Do not make to modifications to EVMS objects unless NCS is running.

---

1. In evmsgui, click *Actions*, select *Create,* then select *Container*.
2. Select *Cluster Segment Manager*, then click *Next*.
3. Select the disks (storage objects) you want to place in the container, then click *Next*.
4. On the Configuration Options page, select the node where you are creating the container, specify *Private* as the type, then choose a name for the container.

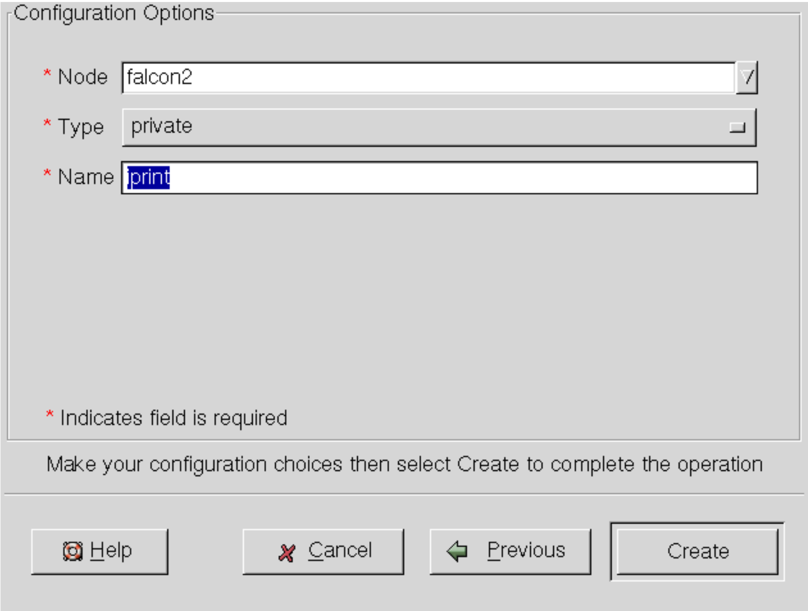

Configuration Options

* Node falcon2

* Type private

* Name print

* Indicates field is required

Make your configuration choices then select Create to complete the operation

The name must be one word, must consist of standard alphanumeric characters, and must not be any of the following reserved words:

- Container
- Disk
- EVMS
- Plugin
- Region
- Segment
- Volume

5. Click *Create* to save your changes.
6. Click OK to close the confirmation dialog and return to the EVMS main screen
7. Click the *Save* icon.

3.9.4 Make New Container Active

1. Click on the *Container* tab.
2. Right-click the new container you just created and select *Modify Properties*.

3. Select your container, and then click *Next*.
4. Click *Modify*.
5. Click the *Save* icon to commit the changes.

3.9.5 Creating an EVMS Volume for iPrint Resources (Conditional)

Verify that EVMS has created your volume by looking at the *Volumes* tab. If the volume does not exist, complete the following:

1. In evmsgui, click *Actions*, select *Create*, and then *EVMS Volume*.
2. Select the container you just created (either the CSM container or the additional segment manager container) and specify a volume name.
3. Click *Create,* then click *Save*.

3.9.6 Creating a File System on the EVMS Volume

You should note the volume name used as it is needed when configuring the cluster scripts.

1. In evmsgui, click the *Volumes* tab and right-click the volume you just created.
2. Select *Make File System*, choose a traditional Linux file system from the list, then click *Next*.

   iPrint requires a native Linux file system such as reiserfs, ext2/3 or xfs. Do not choose Novell NSS or swap.

   You should note the file system used as it is needed when configuring the cluster scripts.
3. Specify a volume label, then click *Make*.
4. Save your changes by clicking *Save*.
5. Exit evmsgui by clicking *Quit* > *Save* > *OK*.

WARNING: Do not mount the file system from EVMSGUI. The file system needs to be mounted by the cluster resource load script. Also, entries should not be found in /etc/fstab file for the operating system to automatically mount this file system.

WARNING: EVMS administration utilities (evms, evmsgui, and evmsn) should not be running when they are not being used. EVMS utilities lock the EVMS engine, which prevents other evms-related actions from being performed. This affects both NSS and traditional Linux volume actions.

NSS and traditional Linux volume cluster resources should not be migrated, offlined, or onlined while any of the EVMS administration utilities are running.

3.10 Configure a Cluster Resource to Host Printing

1 From iManager, click Clusters > Cluster Options.

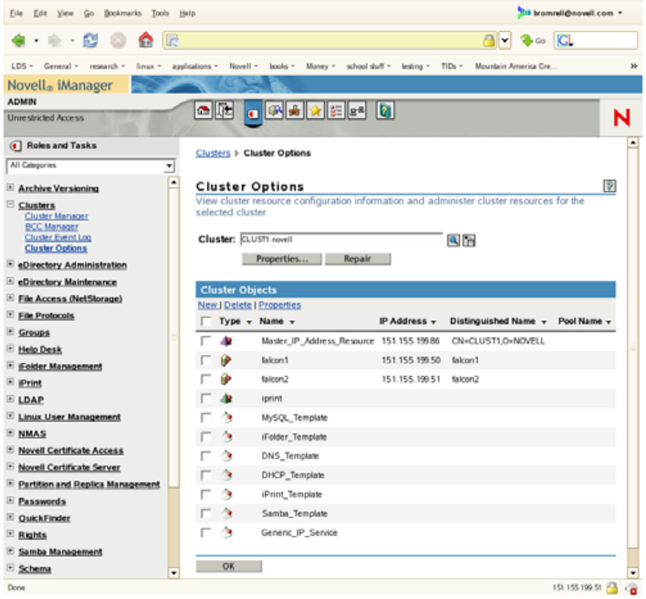

Clusters > Cluster Options

Cluster Options

View cluster resource configuration information and administer cluster resources for the selected cluster.

Cluster: CLUST1.novell

Properties... | Repair

Cluster Objects

New | Delete | Properties

☐	Type	Name	IP Address	Distinguished Name	Pool Name
☐		Master_IP_Address_Resource	151.155.199.86	CN=CLUST1,O=NOVELL	
☐		falcon1	151.155.199.50	falcon1	
☐		falcon2	151.155.199.51	falcon2	
☐		iprint			
☐		MySQL_Template			
☐		iFolder_Template			
☐		DNS_Template			
☐		DHCP_Template			
☐		iPrint_Template			
☐		Samba_Template			
☐		Generic_IP_Service			

OK

2 Select the Cluster Object created during the Section 3.8, "Configuring Novell Cluster Services," on page 23.

3 Under the Cluster Objects heading, click New.

4 Select Resource, and then click Next.

5 Enter a name in the Cluster Resource Name.

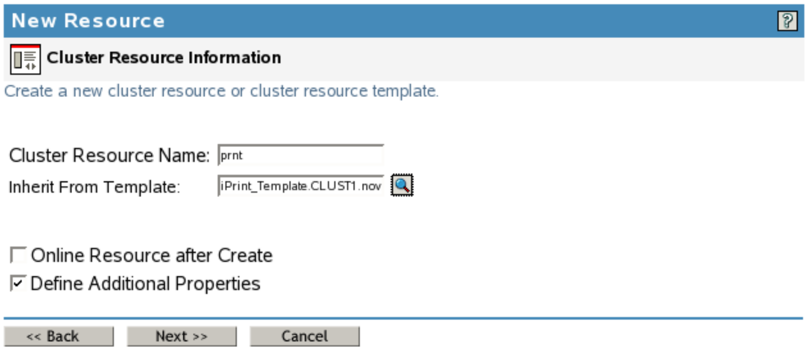

New Resource

Cluster Resource Information

Create a new cluster resource or cluster resource template.

Cluster Resource Name: prnt

Inherit From Template: iPrint_Template.CLUST1.nov

☐ Online Resource after Create

☑ Define Additional Properties

6 Use the object selector for *Inherit from Template* to browse to and open the Cluster object you created, and select the iPrint_Template.

7 Select Define Additional Properties check box, and then click Next.

8 Edit the load script by making the changes outlined in the following table:

Table 3-1 *Load Script Template Changes*

Heading	Description
define the IP address	Replace *a.b.c.d* with the IP address of the iPrint service that is to be used for this resource
define the file system type	Replace *reiserfs* with the correct file system type.
define the container name	Replace *name* with the name of the EVMS container created in section Section 3.9.3, "Creating a Cluster Segment Manager Container," on page 27
define the volume name	Replace *name* with the name of the EVMS volume created in Section 3.9.5, "Creating an EVMS Volume for iPrint Resources (Conditional)," on page 29
define the mount point	Confirm that *mnt/iprint* is the correct mount point, if different mount point for the shared disk other than **/mnt/iprint** is desired. This same mount point must be specified in the unload script.

9 Change the Timeout value to 3 minutes, and then click Next.

10 Edit unload scripts by making the changes outlined in the following table:

Table 3-2 *Unload Script Template Changes*

Heading	Description
define the IP address	Replace *a.b.c.d* with the IP address that is to be used for this resource. This is the same IP address specified in the load script.
define the container name	Replace *name* with the name of the EVMS container. This is the same container name specified in the load script.
define the mount point	Confirm that *mnt/iprint* is the correct mount point, if different mount point for the shared disk other than **/mnt/iprint** is desired. This is the same mount point specified in the load script.

11 Change the Timeout value to 3 minutes, and then click Next.

12 Verify that *Ignore Quorum* is deselected on the Resource Policies page, and then click Next.

13 View the current Preferred Nodes assignments, and if necessary, click the right or left arrow button to assign or unassign servers to the print volume resource.

You can click the up or down arrow button to change the failover order of the servers assigned to the print volume.

If you add a node later to the cluster, it might be added to the Assigned Nodes List. For this reason you should check this list after adding nodes to ensure iPrint fails over to the desired nodes.

14 Click Finish.

3.11 Starting iPrint Cluster Resource

1 In iManager, click Cluster > Cluster Manager.

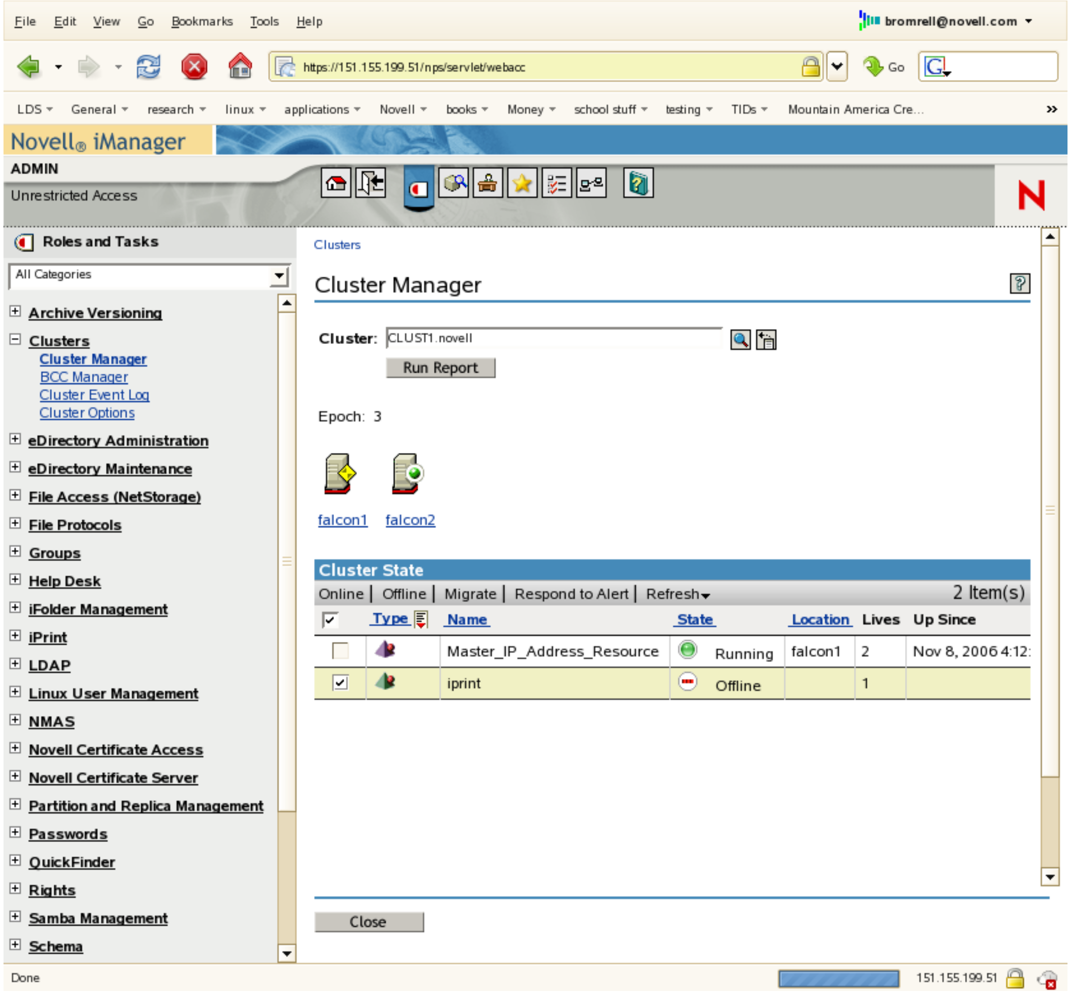

Cluster: CLUST1.novell

Epoch: 3

☑	Type	Name	State	Location	Lives	Up Since
☐		Master_IP_Address_Resource	Running	falcon1	2	Nov 8, 2006 4:12:
☑		iprint	Offline		1	

2 Browse to and select the print Cluster object.

3 Under the Cluster State heading, select the cluster resource you created for print.

4 Click Online to load the iPrint cluster resource.

If the resource does not load, check to ensure that the shared disk is running and that the load script contains the correct entries.

5 Select the node you want to load the resource on.

Now you are ready to begin Creating and Configuring a Driver Store and Print Manager.

3.12 Creating and Configuring a Driver Store and Print Manager

Before running iPrint in a cluster environment, you must create a Driver Store and Print Manager on the cluster-enabled volume.

To create a Driver Store and Print Manager, use iManager.

To create a Driver Store using iManager:

1. From iManager, click iPrint > Create Driver Store.
2. Fill in the fields.

   For more information, see "Creating a Driver Store" in the *OES iPrint Administration Guide for Linux*.

To create a Print Manager:

1. From iManager, click iPrint > Create Print Manager.
2. Fill in the fields.

   For the iPrint service you should use a DNS name that resolves to the secondary IP address used in the cluster load script.
3. When you create the Print Manager for the first time, you might receive an error about trusting a certificate. You need to click on the iPrint Certificate Manager link and accept the certificate, then attempt to create the Print Manager again.

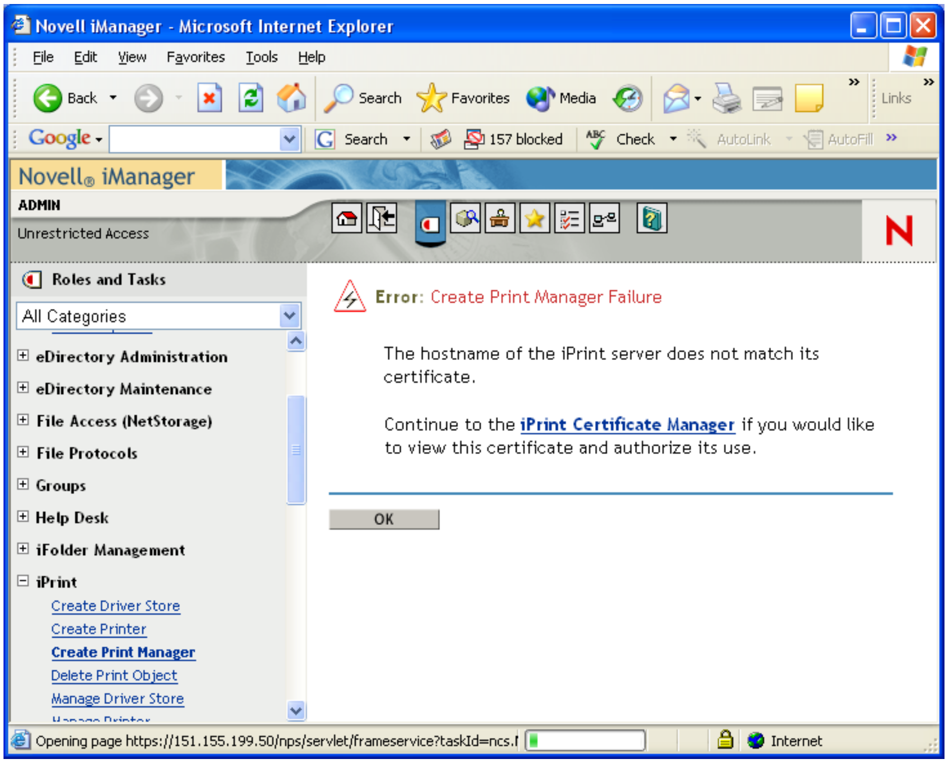

Error: Create Print Manager Failure

The hostname of the iPrint server does not match its certificate.

Continue to the iPrint Certificate Manager if you would like to view this certificate and authorize its use.

For more information, see"Creating a Print Manager" in the *OES iPrint Administration Guide for Linux*.

IMPORTANT: While you can have more than one Print Manager running in a cluster, only one Print Manager can run on a cluster resource node. For example, if you have a four node cluster with serverA, serverB, serverC, and serverD, you can have two print managers: one running on serverA and one on serverB and then assign them to failover to different servers: serverC and serverD. If you assign the print managers to failover to the same cluster resource node, the Print Managers might collide and not failover.

Now you can add printer drivers to the Driver Store and begin creating printers. For more information, see "Adding Printer Drivers" and "Creating a Printer" in the *OES iPrint Administration Guide for Linux*.

3.13 Changing Start, Failover, and Failback Modes for Printing

Novell Cluster Services by default sets the print cluster volume Start Mode and Failover Mode to Auto, and the Failback Mode to Disable. If the print volume Start Mode is set to Auto, iPrint automatically starts on a server when the cluster is first brought up. If the print volume Start Mode is set to Manual, you must manually start print on a server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the print volume Failover Mode is set to Auto, iPrint automatically starts on the next server in the Assigned Nodes list in the event of a hardware or software failure. If the print volume Failover Mode is set to Manual, you can intervene after a failure occurs and before iPrint is moved to another node.

If the print volume Failback Mode is set to Disable, iPrint will not fail back to its most preferred node when the most preferred node rejoins the cluster. If the print volume Failback Mode is set to Auto, iPrint will automatically fail back to its most preferred node when the most preferred node rejoins the cluster. Generally, you want a failback to a preferred node to occur after business hours.You can set the print volume Failback Mode to Manual to prevent iPrint from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

IMPORTANT: We recommend you use the default settings for Start, Failover and Failback modes.

To view or change the print volume Start, Failover, and Failback modes:

1. In iManager, click Clusters > Cluster Options the print volume resource object and select Properties.
2. Browse to and select the Cluster you want to manage.
3. Select the check box of the Print Cluster object you want to modify, and then click the Properties link under Cluster Objects heading.

   The Failover and Failback modes are displayed.

3.14 Using iPrint in a Clustered Environment

If you are using the default iPrint Web pages or are embedding URLs in custom iPrint Web pages, you need to be aware of the following:

- Links to the default iPrint Printer List Web page or other custom pages in \ippdocs should use the DNS name or IP address of the print service. This is the IP Address used in the cluster load script for iPrint. This address moves with the shared volume to other servers in the cluster during a failover.

If you use the server's DNS Name or IP Address or any other IP Address other than the one assigned to the iPrint Resource, links to iPrint Web pages will be broken when iPrint fails over to a different node in the cluster. In this case iPrint links will display a blank page or return an error.

3.15 Troubleshooting iPrint in a Cluster

Use the following sections to troubleshoot common issues when clustering iPrint:

- Section 3.15.1, "Error Code 506D0203," on page 35
- Section 3.15.2, "Resource Goes Comatose on Failover," on page 35
- Section 3.15.3, "iPrint Web Pages Are Displayed on One Node Only," on page 35

3.15.1 Error Code 506D0203

If you receive error 506D0203 and the Print Manager fails to load, ensure you are loading print on a native Linux file system such as ext3. iPrint cannot run on an NSS file system

3.15.2 Resource Goes Comatose on Failover

If a resource goes comatose after a failover, check the nodes in the cluster to see if the shared disk for iPrint is already mounted. If it is mounted, manually unmount the volume using the following command:

```
umount iPrint_mount_point
```

If a message you see a "File system is in use" error, then a user has a terminal open to that file system. You need to have the user, usually an administrator, exit the file system, and then unmount the file system before the cluster resource can come online.

3.15.3 iPrint Web Pages Are Displayed on One Node Only

Ensure the Print Manager is using the secondary IP address for the iPrint resources in /etc/opt/novell/iprint/ipsmd.conf. If the Print Manager was created using the node IP address, then iPrint works only on that node.

3.16 iPrint Configuration Worksheet

To properly configure your system, you need the following setup and information.

- SBD partition of 10 MB.
- Resource Partition on your shared disk of 20 GB.

Use the following worksheet to record the names and IP addresses you need to set up iPrint in a cluster.

Table 3-3 iPrint Cluster Configuration Worksheet

Configuration Information	Where You Need It
Device path for SBD partition:	◆ See Step 2 in Section 3.8, “Configuring Novell Cluster Services,” on page 23.
IP address for the cluster:	◆ See Step 4 in Section 3.8, “Configuring Novell Cluster Services,” on page 23.
Cluster Segment Container Name:	◆ See Step 4 in Section 3.9.3, “Creating a Cluster Segment Manager Container,” on page 27. ◆ See Step 8 in Section 3.10, “Configure a Cluster Resource to Host Printing,” on page 30. ◆ See Step 10 in Section 3.10, “Configure a Cluster Resource to Host Printing,” on page 30.
EVMS Volume Name:	◆ See Step 2 in Section 3.9.5, “Creating an EVMS Volume for iPrint Resources (Conditional),” on page 29. ◆ See Step 8 in Section 3.10, “Configure a Cluster Resource to Host Printing,” on page 30.
Native Linux File System:	◆ See Step 2 in Section 3.9.6, “Creating a File System on the EVMS Volume,” on page 29. ◆ See Step 8 in Section 3.10, “Configure a Cluster Resource to Host Printing,” on page 30.
IP address for the iPrint resource in the cluster:	◆ See Step 8 in Section 3.10, “Configure a Cluster Resource to Host Printing,” on page 30. ◆ See Step 10 in Section 3.10, “Configure a Cluster Resource to Host Printing,” on page 30.

Configuring Samba with Novell Cluster Services

4

4.1 Benefits of Configuring Samba for High Availability

Samba is a suite of tools that let you use the Microsoft* SMB/CIFS networking protocol with Linux computers and other platforms. With Samba installed on a Linux server, client computers can access and use files on the Linux server and shared storage devices connected to the Linux server. If for some reason the Novell® OES Linux server running Samba is not accessible, clients lose their ability to access network files. Configuring Samba with Novell Cluster Services™ ensures that Samba is highly available and network files are accessible. This is possible because the Samba server is automatically started, stopped, and restarted on different servers in the cluster by Novell Cluster Services.

4.2 Samba Installation and Configuration

Novell Cluster Services (NCS) 1.8 must be installed before configuring Samba to work in a cluster. NCS 1.8 provides a Samba resource template, which facilitates configuring Samba in a cluster environment. Novell OES includes NCS 1.8 and licenses for two cluster nodes.

Before running Samba with Novell Cluster Services, Samba must be installed on every server that will run it. To do this, install the Samba server package that ships with Novell OES Linux. This includes the SMB/CIFS file server. You can also install any other desired Samba components. This server package can be installed when you install your OES server or afterward. You can install the server package after the OES installation using either YaST or Linux commands. More information on Samba configuration is available in the /usr/share/doc/packages/samba directory.

4.2.1 Configuring Samba on the Shared Disk

To configure Samba to use shared storage, you need to create a shared directory (file system or disk) on the shared disk system and create mount points to that shared file system on each cluster server that will run Samba.

1. Create a file system on the shared disk system using EVMS.

   Enter `evmsgui` at the Linux server to start the EVMSGUI utility.

   Click Help and use the Web links provided for documentation and information on using EVMS to configure and manage storage.
2. On each cluster node that will run Samba, create the mount point.

   For example, the mount point could be /mnt/samba.
3. Log in as the root user and mount the shared disk (file system) that was created in Step 1.

For example, depending on the mount point and directory names, you could enter a command similar to the following to mount the shared disk:

```
mount /dev/evms/samba /mnt/samba
```

4. At the root of the shared disk you just created (/mnt/samba), enter the following command to create the directories specified:

   ```
   mkdir -p etc/samba
   ```

   These directories must be owned by the root user, and the default group must be root. Also, the directories must have permissions of drwxr-xr-x.

4.2.2 Configuring the smb.conf File

1. Copy the smb.conf file from the /etc/samba directory on one of the OES Linux cluster servers to the etc/samba directory you created on the shared disk in Step 4 on page 10.

   This would be the /mnt/samba/etc/samba directory if you used the same directory names as those given in the example above.
2. Rename the smb.conf file to include the name you will give the Samba cluster resource.

   For example, if the name you will give the Samba cluster resource is samba_resource, you should name the file samba_resource_smb.conf.
3. Configure the Samba daemon by modifying the .conf that you copied in Step 1 and renamed in Step 2.

   3a Add the following lines to the [global] section in the .conf file.

   ```
   bind interfaces only = yes
   interfaces = resource_ipaddress
   pid directory = $SANBA_ROOT/share/locks
   ```

   Replace *resource_ipaddress* with the IP address you plan to assign to the Samba cluster resource.

   3b In the line netbios name = %h-W, change *%h-W* to something unique, such as the name you will give the Samba virtual server.
4. (Conditional) You probably have other instances of Samba running on servers in your cluster. If this is true, edit the smb.conf file on each server where another Samba instance is running and add the following lines to the [global] section:

   ```
   bind interfaces only = yes
   interfaces = server_ipaddress
   ```

   Replace *server_ipaddress* with the IP address of the server where the instance of Samba is running.

   Adding these lines to the smb.conf file eliminates conflicts caused by running multiple instances of Samba.

4.3 Novell Cluster Services Configuration and Setup

After Samba is properly installed and configured , you must create and configure a Samba cluster resource in Novell Cluster Services. This includes configuring Samba load and unload scripts;

setting Samba start, failover, and failback modes; and assigning the Samba resource to specific servers in your cluster.

4.3.1 Creating a Samba Cluster Resource

Novell Cluster Services includes a Samba resource template, which greatly simplifies the process for creating a Samba cluster resource. Much of the Samba cluster resource configuration is perfomed automatically by the Samba resource template. See "Creating Cluster Resources" in the *OES Novell Cluster Services 1.8.1 Administration Guide for Linux* for more information.

To create a Samba cluster resource:

1. Ensure the shared disk (file system) you created and mounted in Step 2 on page 9 is unmounted.

   If you used the directory names specified in the example, you can enter `umount /mnt/samba` to unmount the shared disk.
2. Open your Internet browser and enter the URL for iManager.

   The URL is http://*server_ip_address*/nps/imanager.html. Replace *server_ip_address* with the IP address or DNS name of a server in the cluster or with the IP address for Apache-based services.
3. Enter your username and password.
4. In the left column, locate Clusters and then click Cluster Options.

   iManager displays three links under Clusters that you can use to configure and manage your cluster.
5. Specify the cluster name, or browse and select it, then click New.
6. Specify Resource as the resource type you want to create by clicking the Resource radio button, then click Next.
7. Specify a name for the Samba resource.
8. Type the Samba template name in the Inherit From Template field, or browse and select it from the list.
9. Check the Define Additional Properties check box, click Next, and then continue with Section 4.3.2, "Configuring Samba Load and Unload Scripts," on page 39.

   The Samba resource template configures the Samba resource by automatically creating Samba load and unload scripts, setting failover and failback modes, and assigning Samba as a resource to all nodes in the cluster.

4.3.2 Configuring Samba Load and Unload Scripts

Load Script Configuration

The Samba load script page should already be displayed. The load script contains commands to start the Samba service. You must customize some commands for your specific Samba configuration.

The load script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# define the mount point for the samba partition.
```

```
SAMBA_ROOT=/mnt/samba
# mount the file system
exit_on_error mount -t reiserfs /dev/evms/vol $SAMBA_ROOT
# add the IP address
exit_on_error add_secondary_ipaddress a.b.c.d
# start Samba
exit_on_error /usr/sbin/nmbd --log-basename=$SAMBA_ROOT/share/log --
configfile=$SAMBA_ROOT/share/lib/smb.conf
exit_on_error /usr/sbin/smbd --log-basename=$SAMBA_ROOT/share/log --
configfile=$SAMBA_ROOT/share/lib/smb.conf
# return status
exit 0
```

To customize the Samba load script for your specific configuration:

1. View and, if necessary, edit the following lines for your specific directory structure, mount point, and file system:

   ```
   SAMBA_ROOT=/mnt/samba
   exit_on_error mount -t reiserfs /dev/evms/vol $SAMBA_ROOT
   ```

2. Add the following line below the line to mount the file system:

   `ResourceName=`*`samba_resource_smb.conf`*

   Replace *samba_resource_smb.conf* with the name you gave the smb.conf file in Step 2 on page 38.

3. Edit the following line to assign a unique IP address to the Samba cluster resource:

   ```
   exit_on_error add_secondary_ipaddress a.b.c.d
   ```

   Replace a.b.c.d with the IP address you want to assign to the Samba cluster resource.

   The IP address for the Samba cluster resource allows clients to reconnect to that address regardless of which server is hosting it.

4. Add $ResourceName to the following lines:

   ```
   exit_on_error /usr/sbin/nmbd --log-basename=$SAMBA_ROOT/share/log
   --configfile=$SAMBA_ROOT/share/lib/smb.conf
   exit_on_error /usr/sbin/smbd --log-basename=$SAMBA_ROOT/share/log
   --configfile=$SAMBA_ROOT/share/lib/smb.conf
   ```

   The lines should now appear as follows:

   ```
   exit_on_error /usr/sbin/nmbd --log-basename=$SAMBA_ROOT/share/log
   --configfile=$SAMBA_ROOT/share/lib/$ResourceName
   exit_on_error /usr/sbin/smbd --log-basename=$SAMBA_ROOT/share/log
   --configfile=$SAMBA_ROOT/share/lib/$ResourceName
   ```

5. Click Next and continue with the unload script configuration.

Unload Script Configuration

The Samba unload script page should now be displayed. The unload script contains commands to stop Samba. You must customize some commands for your specific Samba configuration.

The unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
```

```
# define the mount point for the samba partition.
SAMBA_ROOT=/mnt/samba
# request Samba stop
ignore_error killproc -p $SAMBA_ROOT/share/locks/nmbd-smb.conf.pid /
usr/sbin/nmbd
ignore_error killproc -p $SAMBA_ROOT/share/locks/smbd-smb.conf.pid /
usr/sbin/smbd
# stop Samba otherwise
sleep 8
ignore_error fuser -k $SAMBA_ROOT
# del the IP address
ignore_error del_secondary_ipaddress a.b.c.d
# umount the file system
exit_on_error umount $SAMBA_ROOT
# return status
exit 0
```

1. View and, if necessary, edit the mount points in the applicable lines of the unload script.
2. Add the following line below the line where the mount point is defined:

   `ResourceName=`*`samba_resource_smb.conf`*

   Replace *samba_resource_smb.conf* with the name you gave the smb.conf file in Step 2 on page 38.
3. Replace a.b.c.d with the same IP address you specified in the load script.
4. Add $ResourceName to the following lines:

   ```
   ignore_error killproc -p $SAMBA_ROOT/share/locks/nmbd-smb.conf.pid
   /usr/sbin/nmbd
   ignore_error killproc -p $SAMBA_ROOT/share/locks/smbd-smb.conf.pid
   /usr/sbin/smbd
   ```

   The lines should now appear as follows:

   ```
   ignore_error killproc -p $SAMBA_ROOT/share/locks/nmbd-
   $ResourceName.pid /usr/sbin/nmbd
   ignore_error killproc -p $SAMBA_ROOT/share/locks/smbd-
   $ResourceName.pid /usr/sbin/smbd
   ```
5. Click Next and continue with Section 4.3.3, “Setting Samba Start, Failover, and Failback Modes,” on page 41.

4.3.3 Setting Samba Start, Failover, and Failback Modes

The page to set Start, Failover, and Failback modes should now be displayed.

The Samba resource template sets the Samba resource Start mode and Failover mode to Auto and the Failback Mode to Disable.

If the Samba resource Start mode is set to Auto, the Samba server automatically loads on a designated server when the cluster is first brought up. If the Samba Start mode is set to Manual, you can manually start Samba on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.

If the Samba Failover mode is set to Auto, the Samba server automatically moves to the next server in the Assigned Nodes list in the event of a hardware or software failure. If the Samba Failover mode is set to Manual, you can intervene after a failure occurs and before the Samba server is started on another node.

If the Samba Failback mode is set to Disable, the Samba server will continue running on the node it has failed to. If the Samba Failback Mode is set to Auto, the Samba server automatically moves back to its preferred node when the preferred node is brought back online. Set the Samba Failback mode to Manual to prevent the Samba server from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

View or change the Samba resource Start, Failover, and Failback modes, then click Next and continue with Section 4.3.4, "View or Edit Samba Resource Server Assignments," on page 42.

4.3.4 View or Edit Samba Resource Server Assignments

The page to view or change Samba resource server assignments should now be displayed. The Samba resource template automatically assigns the Samba resource to all nodes in the cluster. The order of assignment is the order the nodes appear in the resource list.

To view or edit Samba resource node assignments or change the server failover order:

1. From the list of unassigned nodes, select the server you want the resource assigned to, then click the Right-arrow button to move the selected server to the Assigned Nodes list.

   Repeat this step for all servers you want assigned to the resource. You can also use the Left-arrow button to unassign servers from the resource.
2. Click the Up-arrow and Down-arrow buttons to change the failover order of the servers assigned to the resource or volume.
3. Click Apply or Finish to save node assignment changes.

4.4 Additional Information

For more information on Samba with OES Linux, see ""Samba Support in OES"" in the *Novell OES SP1 Planning and Implementation Guide*.